Investigating the Impacts of Electric Propulsion on Constellation Replenishment Using a Many-Objective Evolutionary Approach

IEPC-2017-127

Presented at the 35th International Electric Propulsion Conference
Georgia Institute of Technology • Atlanta, Georgia • USA
October 8 – 12, 2017

Lake A. Singh[1], Jose J. Guzman[2], Christopher G. Lorenz[3], Stephanie L. Martin[3]
The Aerospace Corporation, Chantilly, Virginia, 20151

Christopher L. Ranieri[4], Demyan V. Lantukh[5], Gregory Fruth[6]
The Aerospace Corporation, El Segundo, California, 90009

Abstract: Value propositions for electric propulsion commonly revolve around system-level metrics such as mass-to-orbit, and the portion of the vehicle mass budget devoted to propulsion. However, electric propulsion also has the potential to provide constellation and program-level benefits. This work studies the impact of electric propulsion on the robustness of a satellite replenishment launch date decision in a simple, zeroth order model. The results demonstrate that the flexibility in arrival date available when using electric propulsion for primary propulsion has the potential to reduce mission gap in the presence of uncertainty. This work additionally demonstrates that electric propulsion can provide robustness against marginally expanded ranges of uncertainty when compared against a chemical propulsion alternative. These results highlight potential, but additional work beyond this first step will be required to build a concrete programmatic value proposition for transitioning towards all-electric systems.

Nomenclature

a_{attain}	=	final semi-major axis attained from trajectory optimization
a_{tgt}	=	target final semi-major axis
e_{attain}	=	final eccentricity attained from trajectory optimization
e_{tgt}	=	target final eccentricity
h_{apo}	=	apogee altitude
h_{per}	=	perigee altitude
i	=	satellite inclination
i_{attain}	=	final inclination attained from trajectory optimization
i_{tgt}	=	target final inclination
I_{sp}	=	specific impulse
M_0	=	mean anomaly at epoch

[1] Engineering Manager, Performance Modeling and Analysis Department, 14745 Lee Road, CH1-510.
[2] Engineering Specialist, Performance Modeling and Analysis Department, 14745 Lee Road, CH1-510.
[3] Member of the Technical Staff, Performance Modeling and Analysis Department, 14745 Lee Road, CH1-510.
[4] Engineering Manager, Flight Mechanics Department, 2310 E. El Segundo Blvd., M4-942.
[5] Senior Member of the Technical Staff, Flight Mechanics Department, 2310 E. El Segundo Blvd., M4-942.
[6] Director, Flight Mechanics Department, 2310 E. El Segundo Blvd., M4-942.

r_E = Earth radius
t_{F0} = nominal legacy satellite failure date
Δt_F = uncertainty in legacy satellite failure date
t_{L0} = nominal new satellite launch date
Δt_L = uncertainty in new satellite launch date
ΔV = delta-vee, change in spacecraft velocity
X_{state} = state error
Ω = right ascension of the ascending node
ω = argument of perigee

I.Introduction

Electric propulsion (EP) enjoys growing adoption in some space market segments, such as commercial satellites in geostationary orbit (GEO). The motivations traditionally given for this trend are that EP can reduce mass-to-orbit and the portion of the vehicle mass budget devoted to propulsion. Beyond this however, EP has the potential to improve resilience and robustness at the constellation and programmatic level. A key example of this is the recovery story for the first Advanced Extremely High Frequency satellite (AEHF-1). Through the incorporation of an EP system, AEHF-1 was able to successfully reach its design orbit despite the failure of its primary chemical propulsion system. Beyond saving a single vehicle, the EP system onboard AEHF-1 positively impacted the entire constellation by preserving capability which would otherwise have been lost.

This paper reports on the first steps in a study to investigate the impacts of EP on robustness and resilience in the context of replenishing and maintaining a satellite capability in GEO. The study explores these impacts by examining the robustness of the replenishment strategy to uncertainties in satellite failure and launch time in the presence of chemical and EP using the many-objective robust decision making (MORDM) framework [1, 2]. The key hypothesis in this work is that the flexibility in transfer time offered by low-thrust EP options provides added adaptability against uncertainties when compared with chemical impulsive transfers.

II.Approach

A. MORDM Analysis Framework and Tools

MORDM represents a union of many-objective decision support using many-objective evolutionary algorithms (MOEAs) and robust decision making (RDM). The many-objective decision support component identifies high performing alternatives and key trades in a many-objective context for complex decision problems through a structured and iterative process of problem formulation, many-objective optimization, visual analytics, and ultimately negotiated selection. Problem formulation consists of identifying decisions stakeholders have available to them, how stakeholders measure value through objectives, how decisions relate to objectives through models, constraints, and sources and ranges of uncertainty for parameters.

Once the problem formulation is implemented into code through the use of models, many-objective optimization occurs with the aid of MOEAs. This study makes use of the Borg many-objective evolutionary algorithm (MOEA) [3]. Borg is a state-of-the-art hyperheuristic which automatically adapts its use of evolutionary operators based on feedback from the search in order to accelerate convergence. The Borg algorithm is implemented into the Genetic Resources for Innovation and Problem Solving (GRIPS) tool: a parallelized many-objective optimization and decision support tool which links MOEAs with detailed spacecraft models through an application programming interface (API) [4-12]. Using Borg and parallel computing resources, GRIPS explores the decision space given in the problem formulation to find and map the Pareto-efficient frontier in the provided objective space.

Once GRIPS and/or Borg maps the efficient frontier, RDM incorporates uncertainty analysis and scenario discovery to identify how the outcomes can vary given uncertainty in exogenous factors. The models used for many-objective optimization can include well-characterized, best-estimate, probability distributions for uncertain parameters so that options can be evaluated against both the nominal state of the world (SOW) and a plausible ensemble of alternate SOWs. RDM subjects the options on the efficient frontier to additional SOWs with expanded uncertainty. Each member SOW in the expanded ensemble applied to the efficient frontier represents a scenario where potential exogenous factors affect outcomes. Thus, the uncertainty analysis step produces a mapping of how alternatives on the efficient frontier respond when subjected to deep uncertainty in exogenous factors. The RDM step is fundamental in engineering projects as the stakeholders start posing "what if?" scenarios that might be outside of the initial set of well-characterized uncertainties and/or outside the scope of the traditional modeling methods used for the initial design.

Scenario discovery interrogates the response to understand where alternatives meet and fail satisficing thresholds of performance in uncertainty-space. Satisficing thresholds are levels of performance elicited from stakeholders which represent satisfactory performance given the adverse influence of exogenous factors. SOWs in the expanded ensemble which violate satisficing thresholds are considered failing. Scenario discovery enables analysts to identify the ranges of uncertain parameters which most often lead to failure using processes such as the Patient Rule Induction Method (PRIM) or Sobol sensitivity analysis [13, 14]. PRIM, developed by Friedman and Fisher in 1998, determines subregions of the input space which result in high or low values, compared to the average over the input domain, for a given target variable. This allows the analyst to determine ranges of certain control parameters on which to optimize values of the target variable. Applications of PRIM are in a variety of fields including financial risk analysis, social sciences, and industrial process control. In this study, the RDM component of the analysis provides insight into how an EP constellation replenishment strategy differs from a chemical propulsion strategy when considering program-level considerations of robustness to uncertain futures.

The preceding analyses produce a wealth of information which analysts explore through the use of interactive visual analytics. By interacting with and visualizing the data, analysts discover insights from the optimization data, highlight key trades and regions of robustness, and communicate these results back to stakeholders. The communication of results back to stakeholders helps them to learn about their complex decision space. It changes their perception of what their decisions are, the ranges of those decisions they have available to them, how they measure value, what their performance thresholds are, and what their constraints are. Thus, MORDM leverages learning feedbacks across the framework to help stakeholders understand their decision space, and ultimately build consensus through negotiated selection in a structured and rigorous way.

B. Model and Notional Problem Definition

This work considers the replenishment of a legacy single-satellite GEO capability with a new vehicle via a simple, zeroth-order scenario. This scenario represents a first step in the modeling and analysis of both well-characterized and deeply uncertain parameters in a replenishment and maintenance study. The new vehicle has a nominal launch date given as t_{L0}, and the legacy vehicle has a nominal design end-of-life (EOL) given as t_{F0}. For this problem, the program manager can control the nominal launch date of the new vehicle, but cannot alter the nominal EOL of the legacy vehicle. The program manager's control over the launch date is not absolute: exogenous factors such as weather, launch vehicle/facility availability, or schedule slips can cause delays which ultimately lead to a slip in the launch date away from the nominal decision made by the program manager. These exogenous factors thus create uncertainty in the actual launch date for the new vehicle. The program manager's knowledge about the failure of the legacy vehicle is also uncertain: it can fail early or it can continue to be a usable resource long after its designed EOL. The capability provided by the systems is assumed to not be additive, such that there is no performance benefit to having an additional vehicle. In this zeroth-order scenario, the single decision to be made by the program manager is when to schedule the nominal launch of the new vehicle so that it arrives to replace the legacy vehicle before it fails, in the presence of uncertainty in when either event will happen, and with a desire to not waste useful life on the legacy vehicle.

The domain of possible nominal launch dates is 1/1/2018 through 1/1/2023, as illustrated in Figure 1. As mentioned previously, the actual launch date has some uncertainty. This notional formulation assumes that the launch date uncertainty is well-characterized as a uniform distribution given over the range $[t_{L0},\ t_{L0} + \Delta t_L]$, where Δt_L is two months. The nominal EOL for the legacy satellite is given as 1/1/2021, with uncertainty in the EOL well-characterized to a uniform distribution given over the range $[t_{F0} - \Delta t_F,\ t_{F0} + \Delta t_F]$ for Δt_F equal to one year. These well-characterized ranges represent confidence intervals which could be identified through historical analysis for a program of record. Each alternative considered by the MOEA is evaluated over an ensemble of 1,000 sets of nominal launch date delay and legacy satellite failure times which each represent a unique SOW. The ensemble used for this evaluation is the same for every alternative launch date. RDM analysis expands the uniform distributions such that Δt_L is six months and Δt_F is three years. These expanded ranges represent potential deep uncertainty in the launch and failure times. Sources of deep uncertainty tend to be exogenous, and outside of the decision maker's ability to measure or predict. Some example sources of deep uncertainty are technology maturation efforts, budgets, or user needs. The objective vector is: 1) minimize the average time across all evaluated SOWs where the legacy vehicle has failed and the new vehicle has not reached the target orbit, and 2) minimize the average time across all evaluated SOWs where the new vehicle has reached the target orbit and the legacy vehicle has not yet failed. These two objectives represent time when there is gap in mission capability, and wasted remaining life on the legacy vehicle, respectively across a range of potential futures. For RDM analysis, a single nominal satisficing performance threshold is applied such that the replenishment strategy is considered a failure in a given SOW if there is a gap in mission of more than two months. This threshold is representative: programs such as those with strict data continuity requirements might have a more stringent threshold, whereas other programs could have a less stringent threshold.

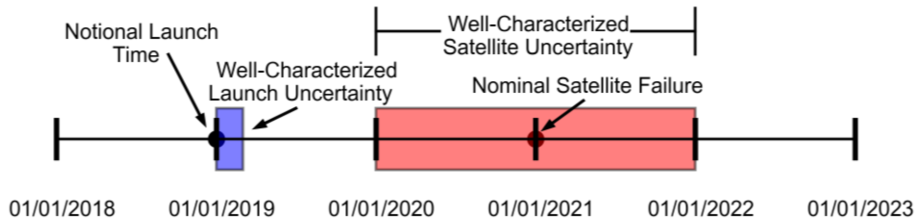

Figure 1. Domain of possible nominal launch dates, with well-characterized uncertainty ranges illustrated.

Table 1 provides the initial and target orbital elements at launch for the new vehicle. The new vehicle begins in a geostationary transfer orbit (GTO), and must transfer to its final orbit with either a chemical or EP system. At the level of fidelity considered for this study, a chemical transfer can be assumed to be impulsive and instantaneous such that the new vehicle arrives at the target orbit on the day of launch. There is consequently no flexibility in when the new vehicle arrives when it uses chemical propulsion. In contrast, if the new vehicle uses EP it does have flexibility in when it arrives. It can follow a minimum propellant trajectory with long time of flight (TOF), a minimum TOF trajectory with high propellant usage, or some trajectory in between. Propellant usage and TOF thus represent two optimization objectives, and a Pareto frontier of alternative transfer strategies exist between them.

Table 1. Initial and target orbital states for the new vehicle.

State	h_{per} (km)	h_{apo} (km)	i	Ω	ω	M_0
Initial	12,821.86	38,421.86	10.0°	300.0°	180.0°	0.0°
Target	35,785.86	35,785.86	0.00°	Any	Any	Any

For the EP case, the new vehicle selects a transfer off of the frontier based on the conditions at launch. If the legacy satellite has already failed at launch, the new vehicle takes a minimum TOF transfer to fill the gap in mission as quickly as possible. Otherwise, the new vehicle takes a minimum propellant transfer to conserve propellant for station-keeping maneuvers. If the launch has been delayed from its nominal date, then the new vehicle takes the minimum propellant transfer which still allows it to arrive at its nominal arrival date had the launch not been delayed. Additionally, the new vehicle can adapt and re-optimize its transfer strategy to a minimum TOF transfer during transit should the legacy satellite fail after launch. This simple behavior captures the new vehicle's flexibility when using EP.

C. Trajectory Problem Formulation and EQLaw Optimization

EP trajectory calculation is accomplished with an enhanced proximity quotient Lyapunov feedback control law (EQLaw), which can reliably generate a feasible trajectory from a starting condition to a subset of target orbital elements [15]. EQLaw is an extension of Q-law and shares the same ability to vary the amount of coasting in a trajectory with an effectivity cutoff parameter [16]. The performance of the trajectory resulting from EQLaw can be modified by adjusting weighting parameters on the targeted orbital elements and some scaling and constraint functions; this improvement is achieved by having GRIPS and Borg iterate on the element target weights.

The specific parameters adjusted by GRIPS in this work are the weightings for semi-major axis, eccentricity, and inclination, as well as the effectivity cutoff parameter. Additional parameters such as the weighting of the other orbital elements are not optimized because these elements are not targeted in the trajectory problem. Integrating a trajectory with EQLaw control provides final element errors to determine transfer feasibility as well as propellant mass and transfer time to gauge solution optimality. The optimizer aggregates the final element errors from each element using Eq. (1) into a state error value to track the feasibility of transfers. This constraint on state error leads to final orbits which are sufficiently accurate for the fidelity of this study, as demonstrated in the results. The transfer time for feasible transfers and associated propellant consumption are the optimization objectives for the current study. Initial mass for the new vehicle is assumed to be 4,000 kg, with the ability to thrust at 0.6 N and 1,700 s I_{sp}. Allowing the optimizer to vary the effectivity cutoff can lead to options with transfer times greater than 6 months. These long TOF transfers are excluded from study in an attempt to account for radiation dose concerns.

$$X_{state} = \sqrt{\left(\frac{a_{attain} - a_{tgt}}{r_E}\right)^2 + \left(e_{attain}^2 - e_{tgt}^2\right)^2 + \left[\tan^2\left(\frac{i_{attain}}{2}\right) - \tan^2\left(\frac{i_{tgt}}{2}\right)\right]^2} \quad (1)$$

III.Results

A. Nominal Transfer Optimization

Figure 2 plots the efficient frontier identified by GRIPS for the nominal transfer optimization. The horizontal axis plots the TOF in days, and the vertical axis plots the transfer ΔV in km/s. Color represents the effectivity cutoff parameter, where a value near one represents a very restrictive thrust timeline and a value of zero represents constant thrusting throughout the transfer. Each circle represents a different transfer profile which can be selected for the new vehicle depending on the conditions at the time of launch. GRIPS identified a number of transfers with high effectivity cutoff parameter which take longer than 180 days to complete. As mentioned previously, this study excludes these options in an attempt to account for radiation dose concerns. Figure 2 shows these options with transparency to indicate that they do not meet this applied TOF constraint.

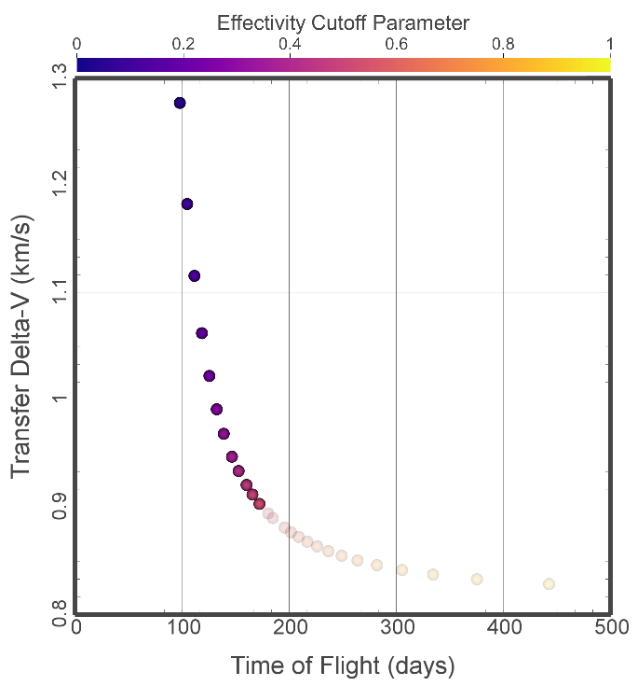

Figure 2. Efficient frontier of transfer options evaluated for time of flight and transfer ΔV. Color represents the EQLaw effectivity cutoff parameter. Transparency is applied to trajectories with TOF greater than 180 days.

Figure 3 shows how the weighting parameters and state error vary across the efficient frontier from Figure 2. Each line in Figure 3 represents a different transfer option represented by a colored dot in Figure 2. As in Figure 2, transparency indicates options which do not meet the 180-day TOF constraint. The results show that the final orbital element errors for the efficient frontier optimized by GRIPS are under 15 km in semi-major axis, 0.0021 for eccentricity, and 0.1° for inclination. Position on the TOF and transfer ΔV front correlates closely with effectivity cutoff, and loosely with the SMA weighting parameter. This indicates a weak relationship between weight application and the resulting trajectory for this problem, or potentially the lack of a unique set of weights to attain an optimal trajectory.

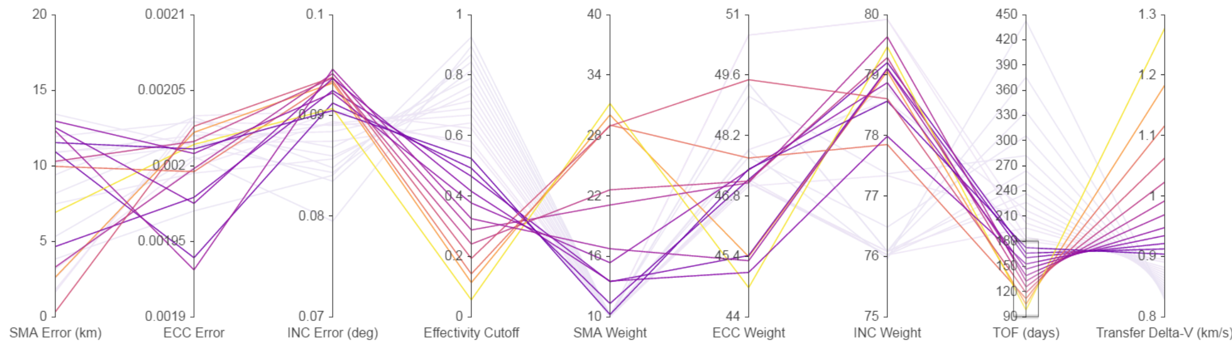

Figure 3. Parallel coordinate plot indicating state error, EQLaw weighting parameters, effectivity cutoff, TOF, and transfer ΔV. Each line represents an option on the efficient frontier from Figure 1. Transparency is applied to trajectories with TOF greater than 180 days.

B. Satellite Replenishment Efficient Frontiers

Figure 4 presents three views of the efficient frontiers generated during optimization of the nominal launch time decision for both EP and chemical systems. Figure 4(A) shows the objective space on the horizontal and vertical axes. To review, the two objectives are to minimize the average times in SOWs generated with well-characterized uncertainty that there are 2 satellites in operation, and no satellites in operation. The population of large orange dots represents the efficient frontier for the EP option, and the small purple dots represent the efficient frontier for the chemical option. Each dot represents a different selection of nominal launch time for either propulsion option. The utopia point of Figure 4(A) is the bottom left of the plot. Thus, Figure 4(A) shows that EP provides a program-level benefit over chemical propulsion under well-characterized uncertainty. Figure 4(B) further demonstrates this by plotting the average time with 0-satellite objective (mission gap) as a function of the nominal arrival date for both options. The EP option can adapt to the realized legacy satellite failure time in different SOWs, such that for a set nominal arrival date the average mission gap in a mission is substantially lowered. When the nominal arrival date is equal to the nominal failure date of 1/1/2021, the difference in average mission gap time is 38 days in the presence of well-characterized uncertainty. This difference becomes more substantial when the nominal arrival date slips past the nominal failure date. This happens because the EP option can effectively arrive early when necessary, and later when able.

The parallel coordinate plot in Figure 4(C) shows each of the parameters on the horizontal and vertical axes of (A) and (B) simultaneously. The range of the 2-satellite objective (mission redundancy) is the same for both propulsion options. This results because, as designed, the EP trajectory selection strategy does not provide a means for adapting to extended mission life on the legacy vehicle. In contrast, the range for the 0-satellite objective is smaller for the EP option. The ability of EP to adapt enables it to reduce the range of this objective value by selecting shorter TOF transfers when appropriate. These results demonstrate that EP can potentially offer a program-level benefit of reducing mission gap in the presence of uncertain launch dates and satellite failure rates.

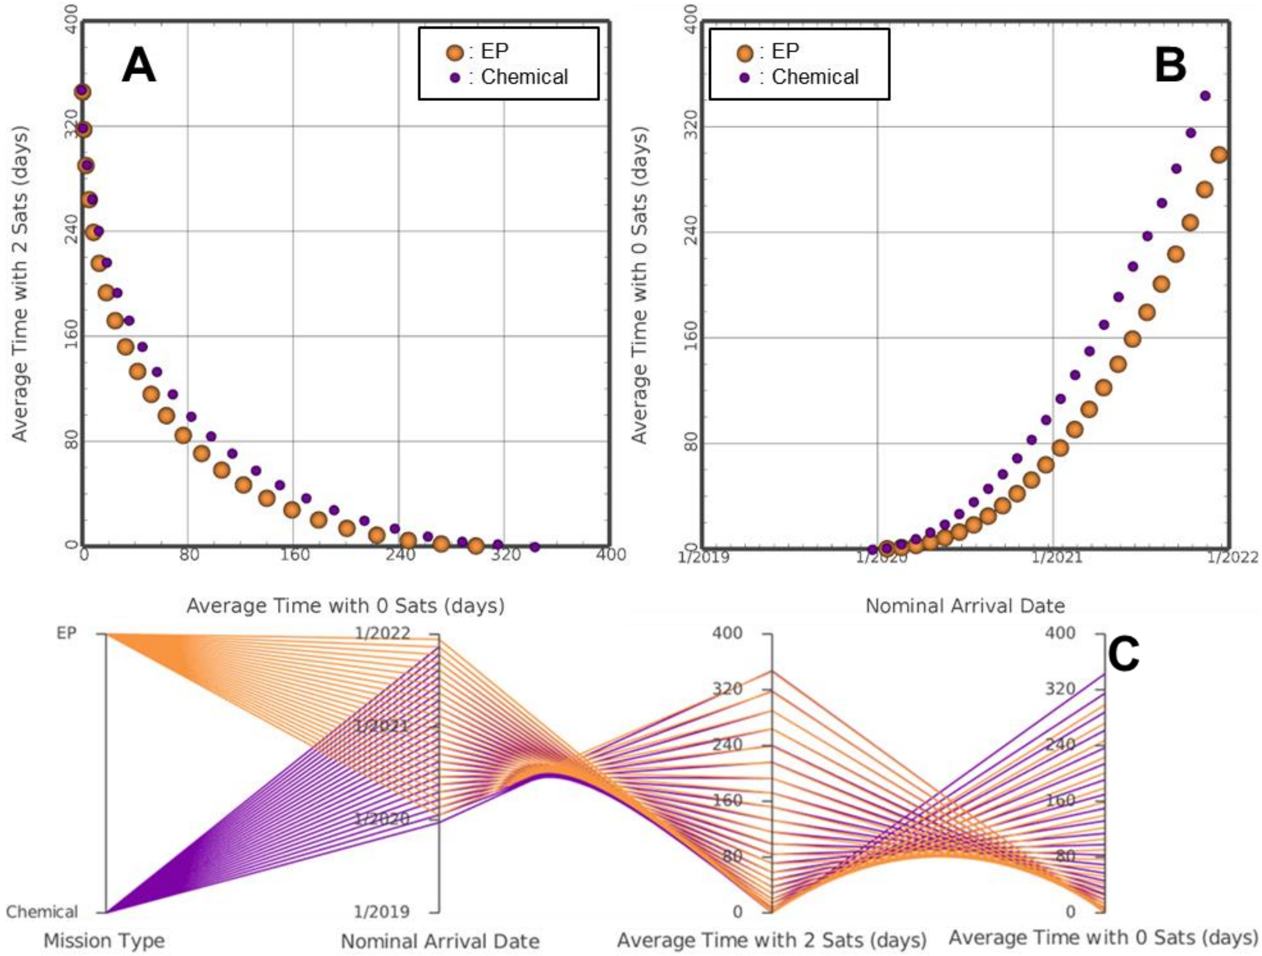

Figure 4. (A) Efficient frontiers for both propulsion options, (B) Average time with 2 satellites in days as a function of the nominal arrival date for both propulsion options, (C) parallel coordinate plot highlighting important parameters.

C. Chemical/Electric Robustness Comparison

The previous section demonstrates the impact of EP on a program-level metric in the presence of well-characterized uncertainty. Figure 5 demonstrates the impact of EP in the presence of an expanded range of deep uncertainty. Figure 5(A) plots the probability of meeting the 2-month mission gap satisficing threshold as a function of the 2-satellite objective. The results demonstrate that, for a given average redundancy in mission, EP offers a higher probability to have less than 2 months of mission gap than chemical propulsion in the presence of deep uncertainty. The maximum difference between EP and chemical options is small, at about 3 percent, but this could vary with additional fidelity. Figure 5(B) similarly highlights the impact of EP on satisficing probability for nominal arrival date. Any given nominal arrival date is slightly "safer" from a satisficing perspective for the EP option. The EP curve is offset from the chemical curve by roughly 3 months, which is the range of adaptability available in arrival time for EP as seen in Figure 2.

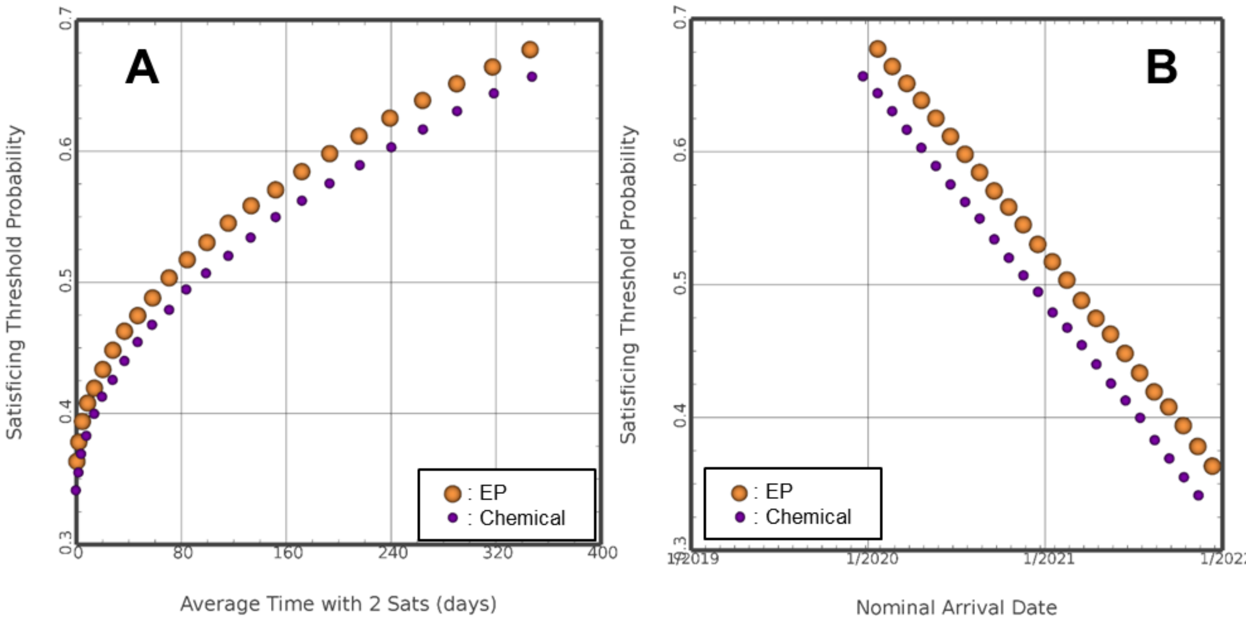

Figure 5. (A) Probability of meeting the 2-month mission gap satisficing threshold as a function of the 2-satellite objective, (B) probability of meeting the 2-month mission gap satisficing threshold as a function of the nominal arrival date.

PRIM analysis reveals that the failure date uncertainty dominates the launch time uncertainty in driving towards failure of the satisficing threshold at all points on either efficient frontier. The effect of either uncertainty source is likely one-for-one, but the deep uncertainty range for the failure date is much larger than the range for the launch date in this work. Figure 6 plots the limit of when the legacy satellite can fail relative to its nominal failure date for the 2-month satisficing threshold to still be attainable as a function of nominal arrival date for both propulsion options. Below each option's line is where the threshold can no longer be met, and above is where it can be. The plot effectively shows the tolerable range of failure date uncertainty given the satisficing threshold and nominal arrival date. The EP line is in general about one month lower than the chemical line. Although the impact is small in this case, it demonstrates that EP can potentially tolerate a wider range of uncertainty than the chemical option because of the ability to adapt.

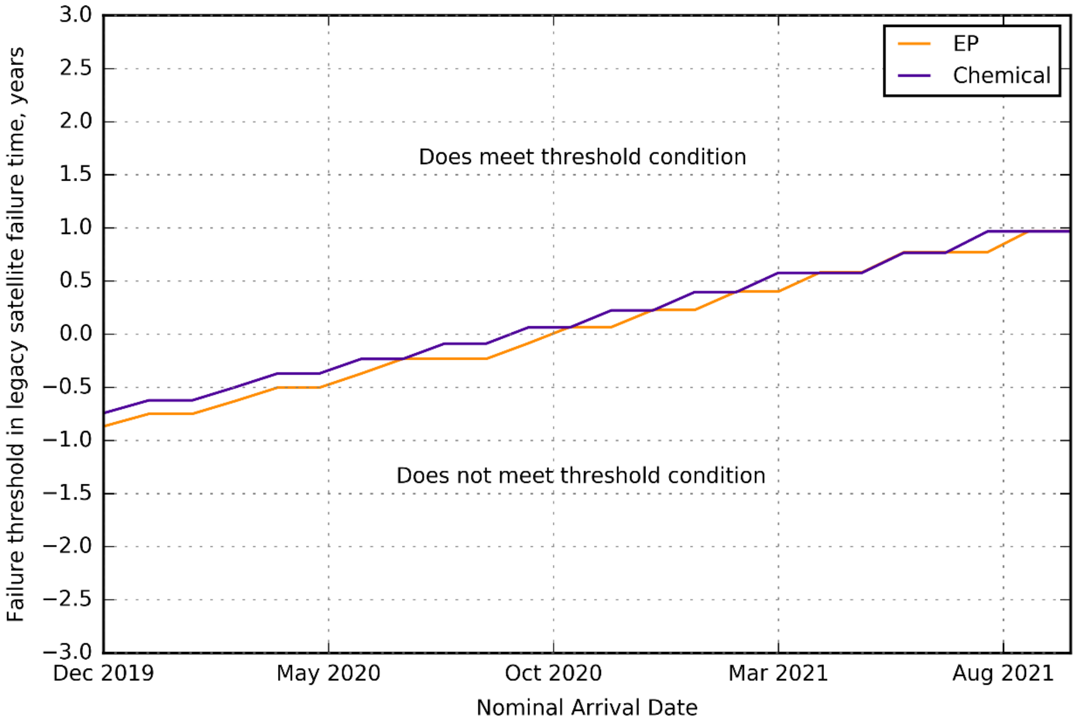

Figure 6. Allowable legacy satellite failure date for the satisficing threshold to still be met as a function of nominal arrival time.

IV.Conclusion

Spacecraft with EP as primary propulsion have flexibility in when they arrive at their target orbit. They can arrive at an earlier time via a minimum TOF transfer, or at a later time via a minimum ΔV transfer. This flexibility has the potential to provide a program-level benefit to program managers as they plan the maintenance of their on-orbit capabilities. This work takes the first steps towards exploring these program-level benefits afforded by EP by studying a simple notional replenishment problem for a one-satellite capability in GEO. By leveraging the EQLaw trajectory optimizer within a MORDM analysis framework, this study examines how EP can impact the robustness of a single nominal launch date decision against uncertainty in actual launch date and legacy capability failure time. The results show that the ability to adapt to uncertain launch and failure dates can provide a concrete benefit to possible mission gap.

EP also provides a modest benefit in the presence of deep uncertainty for the simple scenario provided. However, it remains to be seen if and to what degree the identified benefits manifest in the presence of additional fidelity, complexity, and considerations. Future work will open more decisions to the program manager such as specific impulse, and explore the impacts of EP in the context of additional programmatic considerations such as mission life.

References

1. Kasprzyk, J. R., Nataraj, S., Reed, P.M., Lempert, R.J. "Many objective robust decision making for complex environmental systems undergoing change," *Environmental Modeling and Software* Vol. 42, 2013, pp. 55-71.
doi: 10.1016/j.envsoft.2012.12.007
2. Lempert, R. J. "A new decision sciences for complex systems.," *Proceedings of the National Academy of Sciences* Vol. 99, 2002, pp. 7309-7313.
3. Hadka, D., Reed, P. "Borg: An Auto-Adaptive Many-Objective Evolutionary Computing Framework," *Evolutionary Computation* Vol. 21, No. 2, 2013, pp. 231-259.
doi: 10.1162/EVCO_a_00075

4. Ferringer, M., Clifton, R., Thompson, T. "Systems and Methods for a Core Management System for Parallel Processing of an Evolutionary Algorithm." U.S. Patent 8,433,662, Apr. 30, 2012.
5. Ferringer, M., Clifton, R., Thompson, T. "Systems and Methods for an Application Program Interface to an Evolutionary Software Program." U.S. Patent 8,504,496, Aug. 6, 2013.
6. Ferringer, M., Clifton, R., Thompson T. "Systems and Methods for Parallel Processing Optimization for an Evolutionary Algorithm." U.S. Patent 8,255,344, Aug. 28, 2012.
7. Ferringer, M., Thompson, T. "Systems and Methods for Parallel Processing with Infeasibility Checking Mechanism." U.S. Patent 8,255,345, Aug. 28, 2012.
8. Ferringer, M., Thompson, T. "Systems and Methods for Generating Feasible Solutions from Two Parents for an Evolutionary Process." U.S. Patent 8,494,988, Jul. 23, 2013.
9. Ferringer, M., Thompson T. "Systems and Methods for a Box Fitness Termination of a Job of an Evolutionary Algorithm." U.S. Patent 8,498,952, Jul. 30, 2013.
10. Ferringer, M., Thompson, T. "Systems and Methods for Generating Random Feasible Solutions for an Evolutionary Process." U.S. Patent 82,85,653, Oct. 9, 2012.
11. Ferringer, M., Thompson, T., Clifton, R., DiPrinzio, M. "Systems and Methods for Supporting Restricted Search in High-Dimensional Spaces." U.S. Patent 8,560,472, Oct. 15, 2013.
12. Ferringer, M. P., Clifton, R.S., Thompson, T.G. "Efficient and Accurate Evolutionary Multi-Objective Optimization Paradigms for Satellite Constellation Design," *Journal of Spacecraft and Rockets* Vol. 44, No. 3, 2007, pp. 682-691.
13. Friedman, J. H., Fisher, N.I. "Bump hunting in high-dimensional data," *Statistics and Computing* Vol. 9, No. 123-143, 1999.
14. Sobol, I. M. "Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimation," *Mathematics and Computers in Simulation* Vol. 55, 2001, pp. 271-280.
15. Lantukh, D. V., Ranieri, C.L., DiPrinzio, M.D., Edelman, P.J. "Enhanced Q-Law Lyapunov Control for Low-Thrust Transfer and Rendezvous Design," *2017 AAS/AIAA Astrodynamics Specialist Conference*. Stevenson, WA, 2017.
16. Petropoulos, A. E. "Low-Thrust Orbit Transfers Using Candidate Lyapunov Functions with a Mechanism for Coasting," *AIAA/AAS Astrodynamics Specialist Conference and Exhibit, 16-19 August 2004*. Providence, RI, 2004.